

Provably Manipulation-Resistant Reputation Systems

Paul Christiano*

Abstract

We consider a community of users who must make periodic decisions about whether to interact with one another. We propose a protocol which allows honest users to reliably interact with each other, while limiting the damage done by each malicious or incompetent user. The worst-case cost per user is sublinear in the average number of interactions per user and is independent of the number of users. Our guarantee holds simultaneously for every group of honest users. For example, multiple groups of users with incompatible tastes or preferences can coexist.

As a motivating example, we consider a game where players have periodic opportunities to do one another favors but minimal ability to determine when a favor was done. In this setting, our protocol achieves nearly optimal collective welfare while remaining resistant to exploitation.

Our results also apply to a collaborative filtering setting where users must make periodic decisions about whether to interact with resources such as movies or restaurants. In this setting, we guarantee that any set of honest users achieves a payoff nearly as good as if they had identified the optimal set of items in advance and then chosen to interact only with resources from that set.

1 Introduction

Trust and reciprocity distinguish repeated interactions within a small village from isolated interactions in a massive online community. In a small group we are willing to trust each other and expect that “what goes around comes around,” but online we are more guarded: we rely on trusted mediators, expect others to do the bare minimum they have promised, and rarely leave ourselves vulnerable to exploitation. Can the advantages of trust within a small group be realized at a global scale, in an environment where repeat interactions are rare and manipulation will be rampant if it is profitable?

We help address this question by showing that a subset of individuals who intend to trust and help each other can reap the benefits even in the presence of a large majority of exploitative manipulators. This result continues to hold when feedback is extremely noisy, individuals engage in only a handful of interactions, and the manipulators coordinate adversarially. Our per-user performance guarantees are completely independent of the size of the community.

We use online learning over positive semidefinite matrices to iteratively adjust a symmetric measure of trust. This symmetric measure is coupled with a non-trivial system for tracking reciprocal

*UC Berkeley. Email: paulfchristiano@eecs.berkeley.edu

obligations. The result is a protocol which guarantees that the honest users perform nearly as well as if they had simply excluded all dishonest users.

1.1 Motivating example

Consider a set of N users with periodic opportunities to do one another *favours*. In each round, two users Alice and Bob are chosen at random; Alice may choose to pay a cost of 1, and if she does then Bob receives an expected benefit of $1 + \varepsilon$.

If the users were maximizing collective welfare then they would always take these opportunities. After kN rounds the average user's payoff would be εk . But such an altruistic community is open to exploitation by defectors, who receive favours but don't reliably pay them back. Is there an altruistic strategy which can achieve high collective welfare but isn't vulnerable to such exploitation?

We assume that a favor confers an *expected* benefit of $1 + \varepsilon$, but that the actual benefit is either 0 or ρ . In this case three distinct outcomes are indistinguishable to most users:

- Alice does Bob a favor, but by chance the benefit is 0.
- Alice claims to have done a favor for Bob, though in fact she did not.
- Bob claims not to have received a favor, though in fact he did.

Suppose some set of αN "honest" users is willing to follow whatever protocol we specify. If they knew each other's identities in advance, then after kN rounds these users could achieve an average benefit of $\alpha \varepsilon k$ by always doing favours for each other and never doing unreciprocated favours for outsiders. We show how to obtain a result nearly this good without knowing the identities of the honest users in advance. Namely, we show that there is a protocol which allows the average honest user a payoff of $\alpha \varepsilon k - \mathcal{O}\left(\frac{\rho k^{2/3}}{\alpha}\right)$ after kN rounds, *regardless* of how the dishonest users behave or how many of them there are.

2 The model

Consider a set of N users and a sequence of rounds, $t = 0, 1, \dots$. In each round, a pair of users $x_0^t, x_1^t \in [N]$ is chosen by nature, along with a pair of payoffs $p_0^t, p_1^t \in [-\rho, \rho]$. All users learn the values x_0^t, x_1^t . The users x_i^t choose whether to interact, picking $s_i^t \in \{0, 1\}$. If either user picks 0 then they don't interact and both receive a payoff of 0. If both users pick 1 then they do interact and x_i receives a payoff p_i^t . Each user learns the payoff they receive but not what payoff their partner received.

For the analysis, we fix a set $H \subset [N]$ of "honest" users who will follow whatever protocol we prescribe and whose welfare we will seek to maximize. Our guarantees will simultaneously hold

for any set H whose members follow our protocol. We will assume that the users have access to unlimited cheap talk (in the form of broadcasts), verifiable identities, and a shared random beacon.¹

The *payoff* of a user $p^{\leq T}(x)$ is the sum of their payoffs in rounds they participated in:

$$p^{\leq T}(x) = \sum_{\substack{i, t \leq T: x_i^t = x \\ s_0^t = s_1^t = 1}} p_i^t.$$

Let $p^{\leq T}(H)$ be the total payoff of the honest users after T rounds, i.e.

$$p^{\leq T}(H) = \sum_{x \in H} p^{\leq T}(x).$$

We define the benchmark $\text{OPT}^{\leq T}(H)$ to be the counterfactual payoff of the honest users if they had always interacted with each other and only with each other:

$$\text{OPT}^{\leq T}(H) = \sum_{\substack{t \leq T \\ x_0^t, x_1^t \in H}} (p_0^t + p_1^t).$$

In some of our results we include the possibility of *transfers*. If transfers are available, player x_i^t chooses a payment $\tau_i^t \geq 0$ to make to their partner at the end of round t . Player x_0^t then receives an additional payoff of $\tau_1^t - \tau_0^t$ and player x_1^t receives an additional payoff of $\tau_0^t - \tau_1^t$. These additional payoffs are included in the sum defining $p^{\leq T}(x)$.

Our framework only considers interactions between users, but by introducing users who always agree to interact and always receive a payoff of 0 we can easily model static *resources* such as movies, news articles, or merchants who don't care whom they interact with. Some of our results concern the *collaborative filtering* setting, in which each interaction involves one resource and one user. For notational convenience, we will assume that the resource is always x_1^t and the user is always x_0^t . In this setting the set H contains both honest users and some set of resources. To obtain the best bounds we will usually imagine that H contains the “optimal” set of items according to the users in H , i.e. all items which yield a positive payoff on average.

We have described the interactions as being proposed “by nature” but these interactions can also be proposed by the users. Our protocol can be used as a black box within a larger system designed to maximize $\text{OPT}^{\leq T}(H)$.²

3 The results

Our main result is the following:

¹The assumptions of verifiable identities, broadcast communication, and a random beacon can be replaced by the use of signature schemes and coin flipping protocols or by the availability of a trusted central party.

²This could involve using machine learning or simple heuristics to predict which interactions will be positive, or identifying items a user wants to buy and proposing a popular merchant who sells them, or any number of alternatives.

It also falls on the system to estimate the opportunity cost of a mediocre interaction, as discussed in Section 3.

Theorem 1. *There is a randomized polynomial time protocol which yields expected performance*

$$p^{\leq T}(H) \geq \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho\sqrt[3]{NT^2}\right)$$

for every set H of honest users, assuming that one of the following conditions is satisfied:

1. *Transfers are available.*
2. *In every round $p_1^t = 0$. (The collaborative filtering setting.)*
3. *In every round with $x_0^t, x_1^t \in H$ the payoffs are symmetric ex ante: $\mathbb{E}[p_0^t] = \mathbb{E}[p_1^t]$, conditioned on x_0^t, x_1^t , and all information available before round t .*

Note that the favor setting described in Section 1.1 satisfies this *ex ante* symmetry condition, provided that Alice has an opportunity to do Bob a favor with the same probability that Bob has an opportunity to do Alice a favor.³

If we hold fixed the ratio $k = T/N$, then the regret is $\mathcal{O}\left(\rho N k^{2/3}\right)$. So the introduction of each dishonest user increases the total regret by $\mathcal{O}\left(\rho k^{2/3}\right)$ over the first k interactions.⁴

If we also hold fixed the ratio $\alpha = |H|/N = \alpha$, then the per-user regret bound is *independent* of N . In massive communities, this requires two interacting users to make informed decisions even if they are connected only by a long and tenuous chain of mutual acquaintances.

We also prove the following lower bound on the achievable regret:

Theorem 8. *For every protocol and every T , there exists a set of payoffs and a set H such that*

$$p^{\leq T}(H) \leq \text{OPT}^{\leq T}(H) - \Omega\left(\rho\sqrt{NT}\right)$$

in expectation, even assuming that $p_0^t = p_1^t$ in every interaction.

That is, over the first k interactions it is necessary for the honest users to collectively pay a cost of $\Omega\left(\rho k^{1/2}\right)$ per dishonest user just to get enough information to statistically identify them.

There is a tension between choosing a smaller set H —for which the optimum payoff may be higher—and choosing a larger set H in order to amortize the loss $\mathcal{O}\left(\rho\sqrt[3]{NT^2}\right)$ over a larger set. For collaborative filtering in particular, recommendations optimized for a smaller group of users will be strictly better. This tension is only in the analysis, however, since the guarantee applies simultaneously for every group H . As T grows and the term $\mathcal{O}\left(\rho\sqrt[3]{NT^2}\right)$ becomes less important, the best

³Some condition of this form is clearly necessary. For example, suppose that Alice has an opportunity to do Bob a favor in *every* round. If Alice does not take these opportunities, then we have failed to maximize collective welfare. But if Alice does take these opportunities, then we have failed to maximize Alice’s welfare. Since we need to satisfy the theorem for both $H = \{\text{Alice}\}$ and $H = \{\text{Alice}, \text{Bob}\}$, either way we are in trouble.

⁴ The addition of an honest user also increases the regret by $\mathcal{O}\left(\rho k^{2/3}\right)$, but typically decreases the *average* regret per honest user. If an α fraction of users are honest, then the average regret per honest user is $\mathcal{O}\left(\frac{\rho k^{2/3}}{\alpha}\right)$.

analysis might obtain for smaller and smaller groups H , and to compete the algorithm must make more and more finely-tuned recommendations.

We note that our model differs from many other approaches to reputation systems or collaborative recommendation; our protocol merely *filters* proposed interactions, rather than recommending optimal interactions. It is up to the users of the system to propose interactions, and also to estimate the opportunity cost of mediocre interactions (we assume that any interaction with a positive payoff was good, so 0 needs to represent the “indifferent” outcome; watching a good movie should have a negative payoff, if I would otherwise watch a great movie).

In some settings this seems like an acceptable model. For example, I might browse an online marketplace, find an item that looks like a good deal, and then use the merchant’s reputation to decide whether to purchase it. Or I might see a movie which looks interesting and use reviews to decide whether to watch it.

In other settings this is a very inappropriate model. For example, I might have the choice between hundreds of merchants offering the same item. In this case I just want to be recommended the best one. Our results would need to be strengthened to be applicable in such settings.

See Section 11 for a more extensive discussion of modeling issues, limitations, and directions for further work. Our protocol makes a qualitatively novel guarantee, but there is ample room for improvement. For example, our algorithm is not distributed, the protocol is not incentive-compatible for honest users, and we only consider sequential interactions.

4 The protocol

Our main algorithmic contribution is an online learning algorithm which could be used by a central planner to decide which pairs of users should interact. The primary difficulty is that the set H is not known in advance and so we need to simultaneously lower bound the total payoff of *every* set H . This is an unusual guarantee, and requires a novel technical approach. This algorithm is outlined in Section 4.1 and formally analyzed in Section 6.

This algorithm can then be used directly to achieve our guarantees when $p_1^t = 0$ or when $p_0^t = p_1^t$. If transfers are available, we reduce to the case $p_0^t = p_1^t$ by using transfers to equalize the two payoffs. These reductions are routine, and are analyzed in Section 7.

The most challenging condition in Theorem 1 is *ex ante* symmetry. To address this case we introduce a digital currency, and use transfers of this currency to equalize payoffs. The primary difficulty is ensuring that there is sufficient liquidity to facilitate transactions between honest users. We propose a robust protocol based on logarithmic utility functions. This system is outlined in Section 4.2 and formally analyzed in Section 8.

4.1 Outline of the central planner's algorithm

Consider the following single-payoff problem faced by a central planner: in each round a pair of individuals $x_0^t, x_1^t \in [N]$ is given, the planner outputs $s^t \in [0, 1]$ representing the probability with which x_0^t and x_1^t should interact, and then nature reveals a payoff p^t for the interaction.

There are two different payoffs associated with a set H in this game, depending on whether we think of the payoff as accruing to player 0 or player 1:

$$p_i^{\leq T}(H) = \sum_{t \leq T: x_i^t \in H} s^t p^t.$$

We can create a symmetric version, in which the payoff is shared equally between the two players, by averaging $p_0^{\leq T}$ and $p_1^{\leq T}$. All of these different payoff notions are associated with the same benchmark:

$$\text{OPT}^{\leq T}(H) = \sum_{t \leq T: x_0^t, x_1^t \in H} p^t.$$

We exhibit an algorithm PLAN satisfying:

Theorem 3. *Suppose that $p^t s^t \in [-\rho, \rho]$ and that $\mathbb{E}[p^t] \in [-\rho, \rho]$ conditioned on x_0^t, x_1^t and the information from all previous rounds. Then for every H, T and i , PLAN achieves payoffs satisfying*

$$p_i^{\leq T}(H) \geq \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right)$$

in expectation.

This protocol can be used almost directly to achieve low regret when $p_1^t = 0$ (by applying the lower bound on $p_0^{\leq T}(H)$) or when $p_0^t = p_1^t$ (by applying the lower bound on $p^{\leq T}(H) = p_0^{\leq T}(H) + p_1^{\leq T}(H)$).

The basic outline of our algorithm PLAN is:

- The planner initializes a matrix P_{ab}^t , indexed by $a, b \in [N]$, with $P_{ab}^0 = 1$.
- In step t , the planner outputs $s^t = P_{x_0^t x_1^t}^t$.
- After step t , the planner uses an online learning algorithm to predict matrices C_{ab}^{t+1} (for *cut*) and J_{ab}^{t+1} (for *join*).
- The planner then updates

$$\begin{aligned} P_{ab}^{t+1} &= P_{ab}^t - C_{ab}^{t+1} P_{ab}^t + J_{ab}^{t+1} (1 - P_{ab}^t) \\ &= (1 - J_{ab}^{t+1} - C_{ab}^{t+1}) P_{ab}^t + J_{ab}^{t+1} \end{aligned}$$

The matrices C_{ab}^{t+1} and J_{ab}^{t+1} are chosen by the online learning algorithm with the goal of maximizing $p^{t+1} s^{t+1}$.

J^t and C^t are predicted by the online local learning algorithm from [4], which achieves payoffs as high as the best fixed matrices of the form

$$J_{ab} = \mathbb{P}(\ell(a) = \ell(b) = 0)$$

$$C_{ab} = \mathbb{P}\left((\ell(a) = 0, \ell(b) = 1) \vee (\ell(a) = -1, \ell(b) = 0)\right)$$

for any distribution over labelings $\ell : [N] \rightarrow \{-1, 0, 1\}$.

The most counterintuitive feature of the algorithm is that the predicted matrices are used to adjust P , rather than to directly determine which users should interact. This allows the protocol to build up more complex patterns of interaction as time goes on, and is needed to minimize regret for *every* set H simultaneously. Other approaches we tried, such as using an online learning algorithm with a richer comparison class, were unable to achieve this guarantee.

The analysis proceeds in several steps:

- The online learning guarantee implies that the matrices P_{ab}^t perform nearly as well as any sequence of the form

$$Q_{ab}^t(J^*, C^*) = (1 - J_{ab}^* - C_{ab}^*)P_{ab}^{t-1} + J_{ab}^*,$$

where J^* and C^* are constant matrices from the comparison class. That is, the recommendations J^t and C^t are nearly optimal within this class.

- The strong smoothness lemma for the online learning algorithm implies that J_{ab}^t and C_{ab}^t change slowly. Using this fact and a non-trivial analysis of our update rule we prove that $\left|P_{ab}^t - P_{ab}^{t-1}\right|$ is small on average. This implies that the matrices Q^t perform nearly as well as the matrices

$$R_{ab}^t(J^*, C^*) = (1 - J_{ab}^* - C_{ab}^*)P_{ab}^t + J_{ab}^*.$$

(The superscript on P has changed from $t-1$ to t ; this difference may appear small, but it leads to a much harder analysis and ultimately to our suboptimal dependence on T .)

- We show that if $p_0^{\leq T}(H) \ll \text{OPT}^{\leq T}(H)$, then there is some pair of matrices J^*, C^* in the comparison class such that the matrices $R^t(J^*, C^*)$ substantially outperform the matrices P^t . Namely, we can take $J_{ab}^* = 1$ for $a, b \in H$, and $C_{ab}^* = 1$ for $a \in H, b \notin H$. Similarly, if $p_1^{\leq T}(H) \ll \text{OPT}^{\leq T}(H)$ then we can take $C_{ab}^* = 1$ for $a \notin H, b \in H$.
- (1) and (2) together imply that the matrices $R^t(J^*, C^*)$ cannot much outperform P^t for any J^*, C^* in the comparison class. Combined with (3), this implies that we can't have $p_i^{\leq T}(H) \ll \text{OPT}^{\leq T}(H)$.

Surprisingly, the bottleneck in the quantitative analysis is actually in step (2);⁵ we take a small learning rate to control the differences $\left|P_{ab}^{t+1} - P_{ab}^t\right|$. This is not just a technical requirement. Because the central planner's updates are both optimal and small, we can infer that there are no big adjustments to be made. This property is what lets us ensure that the central planner is making adjustments often enough to “exhaust” all of the good ones.

⁵though if this bottleneck were addressed, the rest of our analysis would need to be tightened in order to actually improve the results

4.2 Protocol for *ex ante* symmetric interactions

Recall the *ex ante* symmetry condition:

$$\mathbb{E} \left[p_0^t - p_1^t \middle| x_0^t, x_1^t, K^{<t} \right] = 0$$

where $K^{<t}$ is a transcript of all of the events prior to round t . In this section we describe a protocol which achieves low regret under this condition.

Our approach is to assign each user a level of wealth $w(x) > 0$. A high value of $w(x)$ indicates a debt owed to a user by the rest of the community, while a low value indicates a debt owed *by* the user. Whenever $p_0^t \neq p_1^t$, we adjust $w(x_0^t)$ and $w(x_1^t)$ so that the “effective payoffs” of the two users are equal and the total wealth $w(x_0^t) + w(x_1^t)$ is conserved. Because the effective payoffs are symmetric, we can apply the central planner’s algorithm to guarantee that the total effective payoff is large. Over the long-run the budget constraints $w(x) > 0$ imply that the actual payoffs can’t be too different from the effective payoffs.

We first considered various straightforward approaches in which payoffs were directly translated into changes in wealth, i.e. which equalized the effective payoffs

$$p_0^t + \Delta w(x_0^t) = p_1^t + \Delta w(x_1^t).$$

The difficulty with these proposals was maintaining the invariant $w(x) > 0$. In order to prevent a violation, we need to prevent interactions with users who have low $w(x)$. But this means that if insufficient liquidity is available, interactions between honest users will stop. This issue made it very difficult to analyze protocols of this type, and in most cases were able to find concrete situations in which these protocols performed poorly,⁶ even in the simple favor setting described in Section 1.1.

Instead we define the effective payoff as the actual payoff plus the change in *log* wealth, and equalize this quantity:

$$p_0^t + \Delta \log \left(w(x_0^t) \right) = p_1^t + \Delta \log \left(w(x_1^t) \right).$$

This automatically leads our learning protocol to place less weight on the payoffs of users who run up substantial deficits, preventing them from causing too much damage. At the same time, we never have to block an interaction between honest users due to insufficient funds; as $w(x) \rightarrow 0$ the system places a higher and higher value on changes in $w(x)$, which prevent it from ever setting $w(x) \leq 0$. The resulting protocol seems to be self-correcting and highly robust to liquidity shortages or imbalances of payments.

⁶Manipulators can break many systems by hoarding currency, either by doing too many nice things or by slightly understating their own payoffs. This creates a liquidity problem amongst the honest users, which is difficult to correct without creating enough currency for the dishonest users to run up destructively large deficits. If this attempted manipulation fails, manipulators can slightly understate their own payoffs when interacting with some honest users and slightly overstate them when interacting with others, creating an imbalance of payments among the honest users which has a similar effect.

5 Related work

In this section we discuss a number of related threads in the literature. Our novel contribution is a performance guarantee in the face of general attacks and in a general environment, rather than considering narrower threat models or making restrictive assumptions about the domain.

Collaborative filtering: There is a large literature on collaborative filtering spanning a wide range of techniques; see [19] for an overview.

Technically the most closely related results are matrix prediction algorithms that achieve low regret compared to the comparison class of matrices with small trace norm, and in particular the recent results for online matrix prediction due to Hazan, Kale, and Shalev-Shwartz [9] and the similar work by the author [4]. However, filtering systems based directly on these algorithms would be vulnerable to manipulation.

Another line of related work investigates the recovery of preference functions; for example see [1, 2, 7]. Some of this work uses assumptions similar to our implicit assumptions, for example the existence of large groups with concordant preferences. This line of work differs in that it aims to reconstruct the entire preference table rather than simply filtering proposed interactions. This stronger goal is unattainable in our noisy and adversarial setting.

Competitive collaborative learning: From a modeling perspective, our work is most similar to, and inspired by, the competitive collaborative learning framework of Awerbuch and Kleinberg [3]. In this setting each user must choose a resource in a series of rounds, and the aim is to ensure that the honest players perform as well as if they had all chosen the single best fixed resource in every round. Like us, they rely on machinery from online learning and convert online learning guarantees into regret bounds for the honest users.

However, this work focuses on identifying an optimal resource to use in every round, rather than considering binary decisions across varying conditions. As a result, we use quite different technical tools and are applicable in very different settings.

Transitive trust: Another literature explicitly examines chains of trust relationships, by which one user can indirectly trust another; for example see [10, 15]. The techniques in this literature tend to focus on a binary notion of links as being either honest or dishonest. This assumption is appropriate in some settings (for example with respect to cryptographic credentials) but not in our target applications.

Manipulation-resistance: Another literature addresses the question of modifying recommendation systems to limit the influence of sybils, fake identities controlled by an attacker. For example, see [22, 14]. The robustness guarantees achieved in this literature are similar to our own, but the results are restricted to the setting where a decision-maker can get useful feedback from the same source across many different decisions. In our setting only a handful of users have information about each decision, and so a decision-maker can only rarely get direct advice from the same source more than once. In exchange, we achieve weaker bounds on the damage done per sybil.

Economic systems for P2P networks: Many decentralized economic protocols have been proposed to reduce the scope for free-riding and enforce reciprocity, especially in the context of Peer-

to-Peer (P2P) networks; for some representative examples see [5, 13, 20, 8, 17, 16].

Driven by our desire to obtain end-to-end performance guarantees, we focus on robustly addressing liquidity constraints and possible payment imbalances, issues which have received little theoretical attention. These issues are particularly important because we assume that users only sometimes have the opportunity to provide valuable services, rather than assuming that such opportunities are always readily available.

Conversely, we do not confront any of the practical issues involved in tracking payments. Our protocol would need to be implemented on top of another system for managing digital payments. We also rely on our symmetric reputation system to provide appropriate incentives for honest reporting, sidestepping another prominent issue in the existing literature.

Other reputation systems: Reputation systems have been the focus of substantial research across a wide range of contexts, for example see [11, 12, 21]. P2P networks and online marketplaces have enjoyed the most attention.

To make a decision in these systems, a user consults feedback from peers who have made similar decisions, and trusts each source of feedback according to their past reliability. This is conceptually similar to our framework, although we use different technical tools and the resulting protocol is much more resistant to manipulation.

It may be possible to bridge the gap between our proposal and more practical reputation systems. For example, EigenTrust [12] is a reputation system which uses random walks to compute the maximum eigenvector of an interaction graph. Although it is vulnerable to manipulation, the protocol is much more efficient than ours. By changing the online learning algorithm used by the central planner in our protocol, it might be possible to replace our use of semidefinite programming by a distributed algorithm based on random walks, achieving the best of both approaches.

References

- [1] Alon, Awerbuch, Azar, and Patt-Shamir. Tell me who I am: An interactive recommendation system. In *SPAA: Annual ACM Symposium on Parallel Algorithms and Architectures*, 2006.
- [2] Awerbuch, Azar, Lotker, Patt-Shamir, and Tuttle. Collaborate with strangers to find own preferences. *MST: Mathematical Systems Theory*, 42, 2008.
- [3] Awerbuch and Kleinberg. Competitive collaborative learning. In *COLT: Proceedings of the Workshop on Computational Learning Theory*, Morgan Kaufmann Publishers, 2005.
- [4] Paul Christiano. Online local learning via semidefinite programming. In *STOC: ACM Symposium on Theory of Computing (STOC)*, 2014.
- [5] B. Cohen. Incentives build robustness in bittorrent. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, USA, 2003.
- [6] Pranav Dandekar, Ashish Goel, Ramesh Govindan, and Ian Post. Liquidity in credit networks: A little trust goes a long way. *CoRR*, abs/1007.0515, 2010.
- [7] Drineas, Kerenidis, and Raghavan. Competitive recommendation systems. In *STOC: ACM Symposium on Theory of Computing (STOC)*, 2002.
- [8] Minaxi Gupta, Paul Judge, and Mostafa Ammar. A reputation system for peer-to-peer networks. In *Proceedings of the 13th International Workshop on Network and Operating System Support for Digital Audio and Video Archive (NOSSDAV-03)*, pages 144–152, New York, June 1–3 2003. ACM Press.
- [9] Elad Hazan, Satyen Kale, and Shai Shalev-Shwartz. Near-optimal algorithms for online matrix prediction. *CoRR*, abs/1204.0136, 2012.
- [10] Jacek Jonczyk and Rolf Haenni. Credential networks: a general model for distributed trust and authenticity management. In *PST*, 2005.
- [11] Audun Jøsang, Roslan Ismail, and Colin Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, 2007.
- [12] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in P2P networks. In Gusztáv Hencsey, Bebo White, Yih-Farn Robin Chen, László Kovács, and Steve Lawrence, editors, *WWW*, pages 640–651. ACM, 2003.
- [13] M. Meulpolder, J. A. Pouwelse, D. H. J. Epema, and H. J. Sips. Bartercast: A practical approach to prevent lazy freeriding in P2P networks. In State University of New York Yuanyuan Yang, editor, *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium*, pages 1–8, Los Alamitos, USA, May 2009. IEEE Computer Society.
- [14] Paul Resnick and Rahul Sami. The influence limiter: provably manipulation-resistant recommender systems. In Joseph A. Konstan, John Riedl, and Barry Smyth, editors, *RecSys*, pages 25–32. ACM, 2007.
- [15] Paul Resnick and Rahul Sami. Sybilproof transitive trust protocols. In John Chuang, Lance Fortnow, and Pearl Pu, editors, *EC*, pages 345–354. ACM, 2009.

- [16] Sven Seuken and David C. Parkes. Sybil-proof accounting mechanisms with transitive trust. In Ana L. C. Bazzan, Michael N. Huhns, Alessio Lomuscio, and Paul Scerri, editors, *AAMAS*, pages 205–212. IFAAMAS/ACM, 2014.
- [17] Sven Seuken, Jie Tang, and David C. Parkes. Accounting mechanisms for distributed work systems. In Maria Fox and David Poole, editors, *AAAI*. AAAI Press, 2010.
- [18] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- [19] Xiaoyuan Su and Taghi M. Khoshgoftaar. A survey of collaborative filtering techniques. *Adv. Artificial Intelligence*, 2009, 2009.
- [20] Vivek Vishnumurthy, Sangeeth Chandrakumar, and Emin Gun Sirer. KARMA : A secure economic framework for peer-to-peer resource sharing, November 13 2003.
- [21] Li Xiong and Ling Liu. Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Trans. Knowl. Data Eng.*, 16(7):843–857, 2004.
- [22] Haifeng Yu, Chenwei Shi, Michael Kaminsky, Phillip B. Gibbons, and Feng Xiao. DSybil: Optimal sybil-resistance for recommendation systems. In *IEEE Symposium on Security and Privacy*, pages 283–298. IEEE Computer Society, 2009.

6 Appendix: The central planner’s algorithm

6.1 The algorithm

We use the algorithm $\mathcal{OLL}_\varepsilon$ for online local learning from [4]. We will apply that algorithm to labelings $\ell : [N] \rightarrow \{-1, 0, 1\}$, i.e. to $3N \times 3N$ matrices with rows and columns indexed by elements of $[N] \times \{-1, 0, 1\}$. We consider a sequence of matrices $E^t \in \mathbb{R}^{3N \times 3N}$. Write $E^{\leq T} = \sum_{t \leq T} E^t$ and similarly for $E^{< T}$. Then $\mathcal{OLL}_\varepsilon(E^{< T})$ is defined as the solution to the following semidefinite program:⁷

$$\begin{aligned} & \underset{\substack{X \in \mathbb{R}^{3N \times 3N} \\ X \succeq 0 \\ 0 \leq X_{(a,i)(b,j)} \leq 1}}{\operatorname{argmax}} \quad \varepsilon \langle X, E^{< T} \rangle + \log \det (X + I) \end{aligned}$$

where $\langle X, E \rangle$ is the matrix inner product $\sum_{a,i,b,j} X_{(a,i)(b,j)} E_{(a,i)(b,j)}$.

For us, the important properties of $\mathcal{OLL}_\varepsilon$ are given by the following results:

Theorem 9. [4] *For any sequence E^t and any $X \succeq 0$ with entries in $[0, 1]$, we have*

$$\sum_{t < T} \left\langle E^t, \mathcal{OLL}_\varepsilon(E^{< t}) \right\rangle \geq \sum_{t < T} \langle E^t, X \rangle - \mathcal{O} \left(\varepsilon \sum_{t < T} |E^t|_1^2 \right) - \mathcal{O}(N\varepsilon^{-1})$$

⁷The algorithm in [4] maximized under some additional equality constraints, e.g. $\sum_i X_{(a,i)(a,i)} = 1$. For our purposes it doesn’t matter whether we include these or not, and for ease of exposition we omit them. It is easy to verify that this leaves the results of [4] unchanged; using these extra inequalities would also leave our results unchanged.

where $|E^t|_1$ is the sum of the absolute values of the entries of E^t .

Lemma 7. [4] For any $E^{<t}$ and any E^t we have

$$\left| \mathcal{OLL}_\varepsilon(E^{<t}) - \mathcal{OLL}_\varepsilon(E^{\leq t}) \right|_\infty \leq \mathcal{O}\left(\varepsilon |E^t|_1\right)$$

Our statement of Theorem 9 and Lemma 7 are slightly different from the statements in [4], so we present proofs in Section 10.

We can now define the central planner's algorithm PLAN_ε . This algorithm is parametrized by a learning rate $\varepsilon < \frac{1}{16}$; in order to achieve the unparametrized version we decrease ε over time. Let p^t be the sparse matrix with $Z_{x_0^t x_1^t}^t = p^t$. Given $\{Z^t\}_{t < T}$, the following algorithm gives the central planner's recommendations in round T :

- Set $P_{ab}^0 = 1$ for $a, b \in [N]$.
- For $t = 0, 1, \dots, T-1$:
 - Define E^t to be the sparse matrix with the following non-zero entries:

$$\begin{aligned} E_{(a,0)(b,0)}^t &= (1 - P_{ab}^t) Z_{ab}^t \\ E_{(a,0)(b,1)}^t &= -P_{ab}^t Z_{ab}^t \\ E_{(a,-1)(b,0)}^t &= -P_{ab}^t Z_{ab}^t. \end{aligned}$$

- Let $X^{t+1} = \mathcal{OLL}_{\varepsilon/\rho}(E^{\leq t})$.
- Define

$$\begin{aligned} C_{ab}^{t+1} &= \frac{1}{4} X_{(a,0)(b,1)}^{t+1} + \frac{1}{4} X_{(a,-1)(b,0)}^{t+1} \\ J_{ab}^{t+1} &= \frac{1}{4} X_{(a,0)(b,0)}^{t+1} + \varepsilon^{1/2} \end{aligned}$$

- Update P by $P_{ab}^{t+1} = (1 - C_{ab}^{t+1} - J_{ab}^{t+1}) P_{ab}^{t+1} + J_{ab}^{t+1}$. The construction of E^t ensures that $p^t P_{x_0^t x_1^t}^t$ is a linear function of $\langle X^t, E^t \rangle$.
- Output $s^T = P_{x_0^T x_1^T}^T$.

6.2 Analysis

For convenience, we will write $p^t(P)$ for the payoff the planner would obtain by using the matrix P in round t , namely

$$p^t(P) = p^t P_{x_0^t x_1^t}.$$

We write $U(P^t, X)$ for the matrix that would be obtained by updating P^t with the matrix $X^{t+1} = X$, i.e. defined by the equations

$$\begin{aligned} C_{ab} &= \frac{1}{4}X_{(a,0)(b,1)} + \frac{1}{4}X_{(a,-1)(b,0)}, \\ J_{ab} &= \frac{1}{4}X_{(a,0)(b,0)} + \varepsilon^{1/2}, \\ U(P^t, X)_{ab} &= (1 - C_{ab} - J_{ab})P_{ab}^t + J_{ab} \end{aligned}$$

Lemma 1. *For any $X \succeq 0$ with $0 \leq X_{(i,a)(j,b)} \leq 1$ we have*

$$\sum_{t \leq T} p^t(P^t) \geq \sum_{t \leq T} p^t(U(P^{t-1}, X)) - \mathcal{O}\left(\frac{\varepsilon}{\rho} \sum |p^t|^2\right) - \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right),$$

where P^t are the matrices computed by PLAN_ε .

Proof. We have $p^t(P^t) = p^t(U(P^{t-1}, X^t))$. Moreover,

$$\begin{aligned} p^t(U(P^{t-1}, X)) &= p^t((1 - J_{x_0^t x_1^t} - C_{x_0^t x_1^t})P_{x_0^t x_1^t}^{t-1} + J_{x_0^t x_1^t}) \\ &= p^t(P_{x_0^t x_1^t}^{t-1}) + J_{x_0^t x_1^t} p^t(1 - P_{x_0^t x_1^t}^{t-1}) - C_{x_0^t x_1^t} p^t P_{x_0^t x_1^t}^{t-1} \\ &= p^t(P^{t-1}) + \frac{1}{4} \langle X, E^t \rangle + \beta^t, \end{aligned}$$

where J and C are defined from X in the same way that J^t and C^t are defined from X^t , and $\beta^t = \varepsilon^{1/2} p^t(1 - P_{x_0^t x_1^t}^{t-1})$ doesn't depend on X . We constructed the matrix E^t precisely so that it would satisfy this identity.

Because $X^t = \mathcal{O}\mathcal{L}\mathcal{L}_{\varepsilon/\rho}(E^{<t})$, by Theorem 9 we have

$$\begin{aligned} \sum_{t \leq T} p^t(P^t) &= \sum_{t \leq T} \left(p^t(P^{t-1}) + \frac{1}{4} \langle X^t, E^t \rangle + \beta^t \right) \\ &\geq \sum_{t \leq T} \left(p^t(P^{t-1}) + \frac{1}{4} \langle X, E^t \rangle + \beta^t \right) - \mathcal{O}\left(\frac{\varepsilon}{\rho} \sum |E^t|_1^2\right) - \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right) \\ &= \sum_{t \leq T} p^t(U(P^{t-1}, X)) - \mathcal{O}\left(\frac{\varepsilon}{\rho} \sum |E^t|_1^2\right) - \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right) \\ &= \sum_{t \leq T} p^t(U(P^{t-1}, X)) - \mathcal{O}\left(\frac{\varepsilon}{\rho} \sum |p^t|^2\right) - \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right), \end{aligned}$$

where the last equality holds because $|E^t|_1 = 3|p^t|$. □

Lemma 2. *For any a, b , we have*

$$\sum_{t < T} |P_{ab}^{t+1} - P_{ab}^t| = \mathcal{O}\left(\frac{\varepsilon^{1/2}}{\rho} \sum_{t < T} |p^t|\right) + \mathcal{O}(1),$$

where P^t are the matrices computed by PLAN_ε .

Proof. Consider the “steady state” value of P given J and C , defined as

$$Y_{ab}^t = \frac{J_{ab}^t}{J_{ab}^t + C_{ab}^t}.$$

This is the value to which P would converge if J and C were left unchanged for many iterations. We will show that Y^t does not change very quickly, and that P^t is “following” Y^t and hence does not change very quickly either.

Note that each entry of J^t and C^t is defined to be a constant multiple of an entry of $\mathcal{OLL}_{\varepsilon/\rho}(E^{<t})$, and that the entry-wise ℓ_1 norm of E^t is $\mathcal{O}(p^t)$. So by Lemma 7, we know that $|J_{ab}^{t+1} - J_{ab}^t| \leq \frac{\varepsilon}{\rho} |p^t|$, and similarly $|C_{ab}^{t+1} - C_{ab}^t| \leq \frac{\varepsilon}{\rho} |p^t|$. By definition we have $C_{ab}^t \geq 0$ and $J_{ab}^t \geq \varepsilon^{1/2}$. Together these facts imply that $|Y_{ab}^{t+1} - Y_{ab}^t| = \mathcal{O}\left(\frac{\varepsilon^{1/2}}{\rho} |p^t|\right)$.

By the definition of P^{t+1} , we have

$$\begin{aligned} P_{ab}^{t+1} &= (1 - J_{ab}^t - C_{ab}^t)P_{ab}^t + J_{ab}^t \\ &= (1 - J_{ab}^t - C_{ab}^t)P_{ab}^t + (J_{ab}^t + C_{ab}^t)Y_{ab}^t \end{aligned}$$

By the definition of J and C , we have $J_{ab}^t + C_{ab}^t \leq \frac{1}{4} + \frac{1}{4} + \frac{1}{4} + \varepsilon^{1/2} \leq 1$. Thus P^{t+1} is a convex combination of P_{ab}^t and Y_{ab}^t . In particular,

$$|P_{ab}^t - Y_{ab}^t| = |P_{ab}^t - P_{ab}^{t+1}| + |P_{ab}^{t+1} - Y_{ab}^t|.$$

Intuitively this suggests that $\sum |P_{ab}^{t+1} - P_{ab}^t| \leq \sum |Y_{ab}^{t+1} - Y_{ab}^t|$, from which the desired result would follow.

To see this formally, consider the total length of the path

$$P_{ab}^0, P_{ab}^1, \dots, P_{ab}^t, Y_{ab}^t,$$

i.e. the sum of the absolute values of the differences between consecutive points. If we change this path by adding a new point:

$$P_{ab}^0, P_{ab}^1, \dots, P_{ab}^t, P_{ab}^{t+1}, Y_{ab}^t$$

we don’t change the total length, because P^{t+1} is in between P^t and Y^t . If we then change this path by changing the last point:

$$P_{ab}^0, P_{ab}^1, \dots, P_{ab}^t, P_{ab}^{t+1}, Y_{ab}^{t+1}$$

we increase the total length by at most $|Y_{ab}^{t+1} - Y_{ab}^t|$.

By induction, the total length of the path is thus $\sum |Y_{ab}^{t+1} - Y_{ab}^t| + \mathcal{O}(1)$, and this total length upper bounds $\sum |P_{ab}^{t+1} - P_{ab}^t|$. \square

The next step is to prove that *if* the honest players are not collectively performing well, *then* there is some recommendation the central planner could have made which would generate a substantial improvement.

The reason for the $\varepsilon^{1/2}$ loss in this theorem is our restriction that $J^t \geq \varepsilon^{1/2}$, which lower bounds the central planner's ability to stop honest and dishonest users from interacting and was needed to prevent rapid changes in P^t .

Lemma 3. *For any set H there is a matrix $X^H \succeq 0$, with $0 \leq X_{(i,a)(j,b)}^H \leq 1$, such that*

$$\text{OPT}^{\leq T}(H) - p_i^{\leq T}(H) = \mathcal{O} \left(\sum_{t \leq T} \left(p^t \left(U \left(P^t, X^H \right) \right) - p^t \left(P^t \right) + \varepsilon^{1/2} |p^t| \right) \right),$$

where P^t are the matrices produced by PLAN_ε .

Proof. Observe that PLAN_ε is perfectly symmetrical under an exchange of x_0^t and x_1^t : the only difference is that the matrices P^t are transposed, and the indices of $X_{(a,i)(b,j)}$ are modified by exchanging $+1$ and -1 . So it suffices to prove the theorem for $i = 0$.

Define the vector v^H via $v_{(a,0)} = 1$ for $a \in H$, $v_{(a,1)} = 1$ for $a \notin H$, and $v_{(a,i)} = 0$ otherwise. Define the matrix X^H as $(v^H)(v^H)^T$. This matrix is clearly PSD, and has entries in $[0, 1]$.

Suppose $a \notin H$. Then $X_{(a,i)(b,j)}^H = 0$, and so $U(P^t, X^H)_{ab} = P_{ab}^t$. Alternatively, suppose $a \in H$. Then $X_{(a,-1)(b,0)} = 0$, and exactly one of $X_{(a,0)(b,0)}$ and $X_{(a,0)(b,1)}$ is 1, according to whether $b \in H$.

So for $a, b \in H$, we have

$$U(P^t, X^H)_{ab} = \frac{3}{4}P_{ab}^t + \frac{1}{4} + \varepsilon^{1/2}(1 - P_{ab}^t).$$

For $a \in H, b \notin H$, we have

$$U(P^t, X^H)_{ab} = \frac{3}{4}P_{ab}^t + \varepsilon^{1/2}(1 - P_{ab}^t).$$

We compute:

$$\begin{aligned} \sum_{t \leq T: x_0^t \in H} \left(U(P^t, X^H) \right) &\geq \frac{3}{4} \sum_{t \leq T: x_0^t \in H} p^t(P^t) + \frac{1}{4} \sum_{t \leq T: x_0^t, x_1^t \in H} p^t - \varepsilon^{1/2} \sum |p^t| \\ &= p_i^{\leq T}(H) - \frac{1}{4} p_i^{\leq T}(H) + \frac{1}{4} \text{OPT}^{\leq T}(H) - \varepsilon^{1/2} \sum |p^t| \end{aligned}$$

Rearranging, and using the fact that $U(P^t, X^H)_{ab} = P_{ab}^t$ for $a \notin H$:

$$\begin{aligned} \text{OPT}^{\leq T}(H) - p_i^{\leq T}(H) &= 4 \sum_{t \leq T: x_0^t \in H} \left(p^t \left(U(P^t, X^H) \right) - p^t(P^t) + \varepsilon^{1/2} |p^t| \right) \\ &= 4 \sum_{t \leq T} \left(p^t \left(U(P^t, X^H) \right) - p^t(P^t) + \varepsilon^{1/2} |p^t| \right), \end{aligned}$$

as desired. \square

We are now ready to prove the main result of this section:

Theorem 2. Suppose that $p^t s^t \in [-\rho, \rho]$ and $\mathbb{E}[p^t] \in [-\rho, \rho]$ conditioned on the preceding rounds. Then

$$p_i^{\leq T}(H) \geq \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho \varepsilon^{1/2} T\right) - \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right)$$

in expectation, where $p_i^{\leq T}(H)$ is the payoff achieved by PLAN_ε .

Proof. By Lemma 3, there exists an $X^H \succeq 0$ with $0 \leq X_{(a,i)(b,j)}^H \leq 1$ such that

$$\begin{aligned} \text{OPT}^{\leq T}(H) - p_i^{\leq T}(H) &= \mathcal{O}\left(\sum_{t \leq T} \left(p^t \left(U(P^t, X^H)\right) - p^t(P^t) + \varepsilon^{1/2} |p^t|\right)\right) \\ &\leq \mathcal{O}\left(\sum_{t \leq T} \left(p^t \left(U(P^t, X^H)\right) - p^t(P^t)\right) + \rho \varepsilon^{1/2} T\right) \end{aligned}$$

in expectation, where the second equality holds because $\mathbb{E}[|p^t|] \leq \rho$. So it suffices to bound $\sum_{t \leq T} \left(p^t \left(U(P^t, X^H)\right) - p^t(P^t)\right)$.

Note that PLAN_ε always satisfies $s^t = P_{x_0^t x_1^t}^t \geq J_{x_0^t x_1^t}^t \geq \varepsilon^{1/2}$, and consequently $|p^t| \leq \frac{s^t}{\rho} = \frac{\varepsilon^{1/2}}{\rho}$.

We have

$$\begin{aligned} &\sum_{t \leq T} \left(p^t \left(U(P^t, X^H)\right) - p^t(P^t)\right) \\ &= \sum_{t \leq T} \left(p^t \left(U(P^t, X^H)\right) - p^t \left(U(P^{t-1}, X^H)\right)\right) + \sum_{t \leq T} \left(p^t \left(U(P^{t-1}, X^H)\right) - p^t(P^t)\right) \\ &\leq \sum_{t \leq T} p^t \left|U(P^t, X^H) - U(P^{t-1}, X^H)\right|_\infty + \mathcal{O}\left(\frac{\varepsilon}{\rho} \sum |p^t|^2\right) + \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right) \\ &\leq \sum_{t \leq T} p^t |P^t - P^{t-1}|_\infty + \mathcal{O}\left(\varepsilon^{1/2} \sum |p^t|\right) + \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right) \\ &\leq \rho \sum_{t \leq T} |P^t - P^{t-1}|_\infty + \mathcal{O}\left(\varepsilon^{1/2} \sum |p^t|\right) + \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right) \\ &\leq \varepsilon^{1/2} \sum_{t \leq T} |p^t| + \mathcal{O}\left(\varepsilon^{1/2} \sum |p^t|\right) + \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right) \\ &= \mathcal{O}\left(\rho \varepsilon^{1/2} T\right) + \mathcal{O}\left(\frac{\rho N}{\varepsilon}\right) \end{aligned}$$

where the third line uses the definition of p^t and Theorem 1, the fourth line uses the definition of $U(\cdot, \cdot)$ and the upper bound $|p^t| \leq \rho \varepsilon^{-1/2}$, the fifth line uses the bound $\mathbb{E}[p^t] \leq \rho$ conditioned on previous rounds (and in particular conditioned on $|P^t - P^{t-1}|_\infty$, which can be computed before round t) the sixth line uses Theorem 2, and the seventh line uses the bound $\mathbb{E}[p^t] \leq \rho$.

□

By setting $\varepsilon = \frac{N^{2/3}}{T^{2/3}}$ and using a standard doubling trick, we obtain a more convenient form:

Theorem 3. *Under the same conditions as Theorem 2, there is an algorithm PLAN for the central planner which achieves an expected payoff of*

$$p_i^{\leq T}(H) \geq \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right).$$

Proof. Every time $T = 2^k$, we start a new copy of PLAN_ε with $\varepsilon = \varepsilon_k = N^{2/3}T^{-2/3}$.

By Theorem 2, we have

$$\begin{aligned} \sum_{\substack{2^k \leq t < 2^{k+1} \\ x_i^t \in H}} p^t s^t &\geq \sum_{\substack{2^k \leq t < 2^{k+1} \\ x_0^t, x_1^t \in H}} p^t - \mathcal{O}\left(\rho \varepsilon_k^{1/2} 2^k\right) - \mathcal{O}\left(\frac{\rho N}{\varepsilon_k}\right) \\ &= \sum_{\substack{2^k \leq t < 2^{k+1} \\ x_0^t, x_1^t \in H}} p^t - \mathcal{O}\left(\rho 2^{2/3k} N^{1/3}\right) \end{aligned}$$

Summing from $k = 0$ to $k = \lceil \log(T) \rceil$, we obtain the desired bound. □

7 Appendix: Algorithms for symmetric payoffs, collaborative filtering, and transfers

7.1 Symmetric payoffs

We say that a game has symmetric payoffs if, in each interaction between honest users, $p_0^t = p_1^t$.

The algorithm SYM followed by an honest user is:

- Initialize a copy of PLAN.
- For $t = 0, 1, \dots$:
 - If you are interacting:
 - * Consult $\text{PLAN}(x_0^t, x_1^t)$ to find the probability s^t with which you should interact.
 - * Using the random beacon decide whether you should interact, and broadcast the result. If you and your partner disagree about whether you should interact, don't interact.
 - * If you decide to interact, observe your payoff and broadcast it. Otherwise broadcast a payoff of 0.

- Report a payoff of $-\rho$ to your instance of PLAN if anything goes wrong:⁸ if the interacting parties broadcast different payoffs, if either party makes multiple broadcasts, if either party broadcasts a payoff outside of $[-\rho, \rho]$, or if you disagree about whether an interaction should have occurred (based on your independent use of the random beacon).
- Otherwise, report a payoff of $p^t = p_0^t/s^t = p_1^t/s^t$ to your instance of PLAN.

We can now apply Theorem 3

Theorem 4. *If $p_0^t = p_1^t$ in any interaction between honest users, then the algorithm SYM achieves a payoff*

$$p^{\leq T}(H) \geq \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho\sqrt[3]{NT^2}\right)$$

Proof. It is easy to verify by induction that all of the honest users' instances of PLAN will be perfectly synchronized, and consequently that nothing will ever go wrong in an interaction between honest users. This also allows us to talk unambiguously about “the” instance of PLAN.

Let p^t be the payoff reported to the central planner in round t . Either $p^t = 0$, $p^t = -\rho$, or $p^t = p_0^t/s^t$ and $p_0^t \in [-\rho, \rho]$. Thus $p^t s^t \in [-\rho, \rho]$; because $p^t = p_0^t/s^t$ with probability at most s^t , and otherwise $p^t \in [-\rho, \rho]$, we have $\mathbb{E}[p^t] \in [-2\rho, 2\rho]$. Thus we can apply Theorem 3 to the payoffs p^t (the difference between ρ and 2ρ makes no difference in the asymptotics).

In any interaction between honest users, we have $\mathbb{E}[p^t] = p_0^t = p_1^t$. Moreover, in any interaction involving an honest user, either their expected payoff is equal to $s^t \mathbb{E}[p^t]$, or else they agreed to interact and $p^t = -\rho$. In either case, their expected payoff is at least $s^t \mathbb{E}[p^t]$.

Thus we can apply Theorem 3 twice and conclude

$$\begin{aligned} \sum_{\substack{i,t:x_i^t \in H \\ s_0^t=s_1^t=1}} p_i^t &\geq \sum_{t:x_0^t \in H} s^t p^t + \sum_{t:x_1^t \in H} s^t p^t \\ &\geq 2 \sum_{t:x_0^t, x_1^t \in H} p^t - \mathcal{O}\left(\rho\sqrt[3]{NT^2}\right) \\ &= \sum_{t:x_0^t, x_1^t \in H} (p_0^t + p_1^t) - \mathcal{O}\left(\rho\sqrt[3]{NT^2}\right) \\ &= \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho\sqrt[3]{NT^2}\right) \end{aligned}$$

in expectation. □

7.2 Collaborative filtering

Our algorithm FILTER for collaborative filtering is a nearly direct application of PLAN. Formally, it works as follows.

⁸In fact outputting a payoff of $-\infty$ would be most natural, since this should never happen in an interaction between honest users, but this would require modifying our analysis.

- Initialize a copy of PLAN.
- For $t = 0, 1, \dots$:
 - If you are x_1^t , output $s_1^t = 1$.
 - If you are x_0^t :
 - * Consult $\text{PLAN}(x_0^t, x_1^t)$ to find the probability s^t with which you should interact.
 - * Using the random beacon, output $s_0^t = 1$ with probability s^t and $s_0^t = 0$ otherwise. Broadcast the result.
 - * If you interact, broadcast your payoff p_0^t . Otherwise broadcast 0.
 - Report a payoff of $-\rho$ to your instance of PLAN if x_0^t broadcasts a payoff outside of $[-\rho, \rho]$ or if you disagree about whether an interaction should have occurred (based on your independent use of the random beacon).
 - Otherwise, report a payoff of $p^t = p_0^t / s^t$ to your instance of PLAN.

Theorem 5. *If $p_1^t = 0$ in any interaction between honest users, then the algorithm FILTER achieves a payoff*

$$p^{\leq T}(H) \geq \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho^3 \sqrt{NT^2}\right)$$

Proof. It is easy to check that the honest user's instantiations of PLAN remain synchronized, that $\mathbb{E}[p^t] \in [-\rho, \rho]$, and that $p^t s^t \in [-\rho, \rho]$. In an interaction where x_0^t is honest, it is easy to see that $\mathbb{E}[s^t p^t] = \mathbb{E}[s_0^t s_1^t p_0^t]$. In an interaction where x_0^t, x_1^t are both honest, then $s_1^t = 1$ and we have $\mathbb{E}[p^t] = p_0^t$.

So we have in expectation

$$\begin{aligned}
 \sum_{\substack{i, t \leq T: x_i^t \in H \\ s_0^t = s_1^t = 1}} p_i^t &= \sum_{t \leq T: x_0^t \in H} s^t p^t \\
 &\geq \sum_{t \leq T: x_0^t, x_1^t \in H} p^t - \mathcal{O}\left(\rho^3 \sqrt{NT^2}\right) \\
 &= \sum_{t \leq T: x_0^t, x_1^t \in H} p_0^t - \mathcal{O}\left(\rho^3 \sqrt{NT^2}\right) \\
 &= \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho^3 \sqrt{NT^2}\right),
 \end{aligned}$$

as desired. □

7.3 Transfers

If the payoffs aren't symmetric, but users are allowed to make transfers, then we can reduce to the symmetric case. That is, given an instance $\mathcal{I}_{\text{transfers}}$ of the general problem with transfers, with payoffs p_i^t , we will produce a new instance $\mathcal{I}_{\text{symmetric}}$ of the symmetric problem, with payoffs q_i^t . We will guarantee that if the users perform well on this synthetic instance $\mathcal{I}_{\text{symmetric}}$, then they also

performed well on $\mathcal{I}_{\text{transfers}}$. Since we know that SYM performs well on $\mathcal{I}_{\text{symmetric}}$, this will yield the desired result.

Whenever an honest user x_i^t has to make a decision in $\mathcal{I}_{\text{transfers}}$ they ask the user x_i^t to make the same decision in $\mathcal{I}_{\text{symmetric}}$. Whenever an honest user is given a payoff p_i^t in $\mathcal{I}_{\text{transfers}}$, they perform the following process in order to generate the payoff q_i^t in $\mathcal{I}_{\text{symmetric}}$:

- Send the message p_i^t to x_{1-i}^t , and receive the message p_{1-i}^t .
- If no message is received, or multiple messages are received, or the message contains a payoff $p_{1-i}^t \notin [-\rho, \rho]$, then set $\tau_i^t = 0$.
- Otherwise set $\tau_i = \max \left\{ \frac{p_{1-i}^t - p_i^t}{2}, 0 \right\}$.
- Report $q_i^t = p_i^t + \tau_{1-i} - \tau_i$.

It is easy to verify that in any interaction between honest users we have

$$q_0^t = q_1^t = \frac{p_0^t + p_1^t}{2}.$$

This allows us to apply the algorithm SYM to the instance $\mathcal{I}_{\text{symmetric}}$.

Theorem 6. *There is an algorithm attaining the payoff*

$$p^{\leq T}(H) \geq \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right)$$

in general if transfers are available.

Proof. We apply the reduction described in this section. In any interaction between honest users we have and in any interaction we have $q_i^t \in [-\rho, \rho]$. So we can apply SYM to the resulting payoffs, and we have in expectation

$$\begin{aligned} p^{\leq T}(H) &= \sum_{\substack{i,t:x_i^t \in H \\ s_0^t = s_1^t = 1}} \left(p_i^t + \tau_{1-i}^t - \tau_i^t \right) \\ &= \sum_{\substack{i,t:x_i^t \in H \\ s_0^t = s_1^t = 1}} q_i^t \\ &\geq \sum_{t:x_0^t, x_1^t \in H} \left(q_0^t + q_1^t \right) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right) \\ &= \sum_{t:x_0^t, x_1^t \in H} \left(p_0^t + p_1^t \right) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right), \\ &= \text{OPT}^{\leq T}(H) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right) \end{aligned}$$

as desired. □

8 Appendix: Ex ante symmetry

As in Section 7.3 we reduce an instance $\mathcal{I}_{\text{ex ante}}$ with ex ante symmetric payoffs to an instance $\mathcal{I}_{\text{symmetric}}$ with symmetric payoffs. At the beginning of the protocol each user initializes a balance $w^0(x) = 1$ for every user x . When x_i^t receives the payoff p_i^t in $\mathcal{I}_{\text{ex ante}}$ they do the following to generate the payoff q_i^t in $\mathcal{I}_{\text{symmetric}}$:

- Communicate the payoff p_i^t to x_{1-i}^t , and hear the payoff p_{1-i}^t .
- If no message is heard, or multiple messages are heard, or the message contains a payoff $p_{1-i}^t \notin [-\rho, \rho]$, then set $\tau_i = 0$
- Otherwise $\tau_i^t = \frac{\delta w_0^t w_1^t (p_{1-i}^t - p_i^t)}{2\rho(w_0^t + w_1^t)}$, where $w_i^t = w^t(x_i^t)$:
- Each user updates $w^{t+1}(x_i^t) = w^t(x_i^t) + \tau_i^t$. For $x \notin \{x_0^t, x_1^t\}$, $w^{t+1}(x) = w^t(x)$.
- Report the payoff $q_i^t = p_i^t + \frac{2\rho}{\delta} \frac{\tau_i^t}{w_i^t}$.

Here $\delta < \frac{1}{2}$ is a parameter of the reduction.

It is easy to verify that in any interaction between honest users, both users will report a payoff of

$$q_0^t = q_1^t = \frac{w_0 p_0^t + w_1 p_1^t}{w_0 + w_1}.$$

This is very similar to the protocol with transfers, except that the payoff of an interaction is now a weighted average of the individual payoffs. Once again, we can apply the algorithm SYM to $\mathcal{I}_{\text{symmetric}}$. The aim the analysis will be to show that if the payoff in $\mathcal{I}_{\text{symmetric}}$ is good, then so is the payoff in $\mathcal{I}_{\text{ex ante}}$, which is less straightforward than in the case with transfers.

It is easy to verify that $|\tau_i^t| \leq \delta w_i^t$, so that each $w^t(x)$ changes by at most a factor of $(1 \pm \delta)$ in any round. Since $\delta < \frac{1}{2}$, this guarantees that $w^t(x) > 0$.

The mechanism used to track the quantities $w^t(x)$ is not particularly important; for simplicity, and because our algorithm is already hopelessly undistributed, we assumed that every user independently tracks the quantities $w^t(x)$ for every other user. It would be straightforward, however, to use a more efficient implementation of digital cash such as [6].⁹ In this case a dishonest player might decline to make an appropriate transfer, but this can be handled in the same way as in Section 7.3.

To analyze this protocol we rely on the potential functions

$$U^T(x) = \frac{2\rho}{\delta} \log(w^T(x)) + p^{<T}(x)$$

We use three lemmas to relate the reported payoffs q_i^t to the actual payoffs p_i^t .

Our first lemma relates the reported payoffs q_i^t to the underlying payoffs p_i^t :

⁹Though this introduces the possibility of attacks on the payment infrastructure itself.

Lemma 4. *Under the ex ante symmetry condition, in any interaction between honest users we have*

$$\mathbb{E} [q_0^t + q_1^t] = \mathbb{E} [p_0^t + p_1^t].$$

Proof. Under the ex ante symmetry condition, conditioned on the values of w_0^t and w_1^t , and on the users' decision to interact, we have $\mathbb{E} [p_0^t] = \mathbb{E} [p_1^t]$. Let $\beta_0 = \frac{w_0^t}{w_0^t + w_1^t}$ and $\beta_1 = \frac{w_1^t}{w_0^t + w_1^t}$. Then we have

$$\begin{aligned} \mathbb{E} [q_0^t + q_1^t] &= 2\mathbb{E} [\beta_0 p_0^t + \beta_1 p_1^t] \\ &= 2\beta_0 \mathbb{E} [p_0^t] + 2\beta_1 \mathbb{E} [p_1^t] \\ &= \beta_0 \mathbb{E} [p_0^t + p_1^t] + \beta_1 \mathbb{E} [p_0^t + p_1^t] \\ &= \mathbb{E} [p_0^t + p_1^t] \end{aligned}$$

as desired. □

Our second lemma relates the reported payoffs q_i^t to the change in $U^t(x)$:

Lemma 5. *If $s_0^t = s_1^t = 1$, then then $U^{t+1}(x_i^t) \geq U^t(x_i^t) + q_i^t - \mathcal{O}(\rho\delta)$.*

Proof. We have

$$\begin{aligned} U^{t+1}(x_i^t) &= \frac{2\rho}{\delta} \log(w^{t+1}(x_i^t)) + p^{< t+1}(x_i^t) \\ &= U^t(x_i^t) + \frac{2\rho}{\delta} \log\left(\frac{w^{t+1}(x_i^t)}{w_i^t}\right) + p^t(x_i^t) \\ &= U^t(x_i^t) + \frac{2\rho}{\delta} \log\left(1 + \frac{\tau_i^t}{w_i^t}\right) + p^t(x_i^t) \\ &> U^t(x_i^t) + \frac{2\rho}{\delta} \frac{\tau_i^t}{w_i^t} - 2\rho\delta + p^t(x_i^t) \\ &= U^t(x_i^t) + q_i^t - 2\rho\delta \end{aligned}$$

where we have used the bound $\left|\frac{\tau_i^t}{w_i^t}\right| \leq \delta$ and the inequality $\log(1+z) > z - \delta^2$ for $|z| \leq \delta \leq \frac{1}{2}$. □

Our third lemma relates the quantities $U^T(x)$ to the payoffs $p^{\leq T}(x)$:

Lemma 6.

$$p^{\leq T}(H) \geq \sum_{x \in H} U^T(x) - \frac{2\rho}{\delta} |H| \log\left(\frac{N}{|H|}\right)$$

Proof. $\sum_x w^t(x)$ is conserved by the algorithm, and $w^t(x) \geq 0$, so

$$\sum_{x \in H} w^T(x) \leq \sum_x w^T(x) = \sum_x w^0(x) = N.$$

By Jensen's inequality, $\sum_x \log(w^T(x)) \leq |H| \log\left(\frac{N}{|H|}\right)$. We have

$$\begin{aligned} p^{\leq T}(H) &= \sum_{x \in H} U^T(x) - \frac{2\rho}{\delta} \sum_{x \in H} \log(w^T(x)) \\ &\geq \sum_{x \in H} U^T(x) - \frac{2\rho}{\delta} |H| \log\left(\frac{N}{|H|}\right), \end{aligned}$$

as desired. □

Combining these three lemmas yields our desired result:

Theorem 7. *There is an algorithm that obtains regret*

$$p^{\leq T}(H) \geq OPT^{\leq T}(H) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right)$$

assuming that the payoffs between honest users satisfy ex ante symmetry.

Proof. We will assume that T is known in advance, and choose $\delta = \sqrt{NT}$; this assumption can then be removed by a standard doubling trick exactly as in Theorem 3.

In any interaction between honest users we have $q_0^t = q_1^t$. Moreover, $q_i^t \in [-\rho, \rho]$. Thus we can apply SYM and Lemma 4 and obtain in expectation

$$\begin{aligned} \sum_{\substack{t, i: x_i^t \in H \\ s_0^t = s_1^t = 1}} q_i^t &\geq \sum_{t: x_0^t, x_1^t \in H} (q_0^t + q_1^t) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right) \\ &= \sum_{t: x_0^t, x_1^t \in H} (p_0^t + p_1^t) - \mathcal{O}\left(\rho \sqrt[3]{NT^2}\right). \end{aligned}$$

Then we can apply Lemma 5 and Lemma 6 and obtain in expectation

$$\begin{aligned}
p^{\leq T}(H) &\geq \sum_{x \in H} U^T(x) - \frac{2\rho}{\delta} |H| \log \left(\frac{N}{|H|} \right) \\
&\geq \sum_{x \in H} U^T(x) - \frac{2\rho N}{\delta} \\
&= \sum_{x \in H} U^T(x) - \mathcal{O}(\rho\sqrt{NT}) \\
&\geq \sum_{x \in H} U^0(x) + \sum_{\substack{i,t:x_i^t \in H \\ s_0^t=s_1^t=1}} q_i^t - \rho\delta T - \mathcal{O}(\rho\sqrt{NT}) \\
&= \sum_{\substack{i,t:x_i^t \in H \\ s_0^t=s_1^t=1}} q_i^t - \mathcal{O}(\rho\sqrt{NT}) \\
&\geq \sum_{t:x_0^t,x_1^t \in H} (p_0^t + p_1^t) - \mathcal{O}(\rho\sqrt[3]{NT^2}) - \mathcal{O}(\rho\sqrt{NT}) \\
&\geq \text{OPT}^{\leq T}(H) - \mathcal{O}(\rho\sqrt[3]{NT^2})
\end{aligned}$$

where the last inequality assumes $T > N$. If $T \leq N$ the result is trivial, since in this case the maximum possible difference between $p^{\leq T}(H)$ and $\text{OPT}^{\leq T}(H)$ is $2\rho T = \mathcal{O}(\rho\sqrt[3]{NT^2})$. \square

9 Appendix: Lower bounds

In this section we consider the game described in Section 6. We show that for any T , there is a set H and a strategy for nature such that the planner's payoff is

$$\text{OPT}^{\leq T}(H) - \Omega(\sqrt{NT})$$

These results extend immediately to all of our other settings, because an algorithm for any of these settings could be trivially used by the central planner (who could simply simulate the interactions of the users and recommend that two users interact with whatever probability they interacted in the simulation).

To prove this result we consider a randomized strategy for nature such that the expected payoff of every user is 0. We then consider the best possible partition of the users into two sets H_1, H_2 , and show that $\text{OPT}^{\leq T}(H_1) + \text{OPT}^{\leq T}(H_2) = \Omega(\sqrt{NT})$. This implies that for some choices by nature, and for one of the sets H_i , the payoff of the central planner is less than $\text{OPT}^{\leq T}(H_i) - \Omega(\sqrt{TN})$.

Theorem 8. *For every algorithm for the central planner and every T , there exists a set of payoffs and a set H such that*

$$\text{OPT}^{\leq T}(H) - o(\sqrt{NT})$$

in expectation.

Proof. Consider a strategy for nature in which each x_0^t and x_1^t is uniformly random, and each p^t is a uniform ± 1 random variable. It is easy to see that no matter how the central planner plays in this game, the expectation of $p^{\leq T}(H) = 0$.

Given some enumeration of the users $0, 1, \dots, N-1$, we define the set H_1 as follows. We include all of the users $0, 1, \dots, N/2$ in H_1 . We then inductively decide whether to include $i > N/2$ in H_1 by evaluating the sum

$$P_x = \sum_{\substack{t \leq T: x_1^t < x_0^t = x \\ x_1^t \in H_1}} p^t + \sum_{\substack{t \leq T: x_0^t < x_1^t = x \\ x_0^t \in H}} p^t.$$

We include x in H_1 if and only if $P_x > 0$.

The sum defining each P_x has in expectation at least $T/(2N)$ terms, because $1/N$ of all interactions involve user x and at least half of these interactions are with a user in $|H_1|$. Each of these terms is a uniform ± 1 random variable. So the probability that $P_x > 0$ is $\frac{1}{2}$ and the expectation of P_x , given that $P_x > 0$, is at least $\sqrt{T/N}$.

It is easy to see that $\text{OPT}^{\leq T}(H_1) = \sum_{x \in H_1} P_x$. Each term P_x for $x \leq N/2$ has expectation 0. In expectation there are $\Omega(N)$ terms with $x > N/2$. Each of these terms is $\Omega(\sqrt{T/N})$ in expectation. Thus $\text{OPT}^{\leq T}(H_1) = \Omega(\sqrt{TN})$.

We define $H_2 = [N] \setminus H_1$. We have $\text{OPT}^{\leq T}(H_2) = 0$: for each pair $x < y$ with $x, y \in H_2$, the event $y \in H_2$ depends only on the sum P_y , which is a sum of terms independent of any interactions occurring between x and y . Thus if there is an interaction between x and y , its expected payoff is 0. Summing over pairs of users in H_2 , we obtain $\mathbb{E}[\text{OPT}(H_2)] = 0$.

Thus $\mathbb{E}[\text{OPT}^{\leq T}(H_1) + \text{OPT}^{\leq T}(H_2)] = \Omega(\sqrt{TN})$, while $\mathbb{E}[p^{\leq T}(H_1) + p^{\leq T}(H_2)] = \mathbb{E}[p^{\leq T}([N])] = 0$. In particular, $\mathbb{E}[\text{OPT}^{\leq T}(H_1) + \text{OPT}^{\leq T}(H_2) - p^{\leq T}(H_1) - p^{\leq T}(H_2)] = \Omega(\sqrt{TN})$, and so we can find some particular strategy by nature and a set H_i such that $p^{\leq T}(H_i) \leq \text{OPT}^{\leq T}(H_i) - \Omega(\sqrt{TN})$, as desired. \square

10 Appendix: Strengthened analysis of online local learning

Lemma 7. [4] For any $E^{<t}$ and any E^t we have

$$\left| \mathcal{OLL}_\varepsilon(E^{<t}) - \mathcal{OLL}_\varepsilon(E^{\leq t}) \right|_\infty \leq \mathcal{O}\left(\varepsilon |E^t|_1\right)$$

Proof. Note that the ℓ_1 and ℓ_∞ norms are dual. So this result follows directly from Lemma 8 in [4] and Lemma 2.19 in [18]. \square

Theorem 9. [4] For any sequence E^t and any $X \succeq 0$ with entries in $[0, 1]$, we have

$$\sum_t \left\langle E^t, \mathcal{OL}\mathcal{L}_\varepsilon(E^{<t}) \right\rangle \geq \sum_t \left\langle E^t, X \right\rangle - \mathcal{O} \left(\varepsilon \sum_{t \leq T} |E^t|_1^2 \right) - \mathcal{O}(N\varepsilon^{-1})$$

where $|E^t|_1$ is the sum of the absolute values of the entries of E^t .

Proof. By Lemma 2.3 of [18] and Lemma 5 of [4], for any $X \succeq 0$ with entries in $[0, 1]$ we have:

$$\sum_t \left\langle E^t, \mathcal{OL}\mathcal{L}_\varepsilon(E^{<t}) \right\rangle \geq \sum_t \left\langle E^t, X \right\rangle - \mathcal{O}(N\varepsilon^{-1}) + \sum_t \left\langle E^t, \mathcal{OL}\mathcal{L}_\varepsilon(E^{<t}) - \mathcal{OL}\mathcal{L}_\varepsilon(E^{\leq t}) \right\rangle.$$

Using Lemma 7 we have

$$\begin{aligned} \sum_t \left\langle E^t, \mathcal{OL}\mathcal{L}_\varepsilon(E^{<t}) - \mathcal{OL}\mathcal{L}_\varepsilon(E^{\leq t}) \right\rangle &\geq - \sum_t |E^t|_1 \left| \mathcal{OL}\mathcal{L}_\varepsilon(E^{<t}) - \mathcal{OL}\mathcal{L}_\varepsilon(E^{\leq t}) \right|_\infty \\ &= \mathcal{O} \left(\sum_t \varepsilon |E^t|_1^2 \right) \end{aligned}$$

which gives the desired result. \square

11 Appendix: Open questions

We have presented a protocol for managing trust which achieves a qualitatively new—and surprisingly strong—formal guarantee. But the theoretical picture is far from complete, and the proposed protocol is far from practical. We hope that our model and algorithmic approach will inspire further work in both of these directions, and in this section we highlight some areas for improvement and elaboration.

Tighter regret bounds: Our regret bound of $\mathcal{O}(\rho^3 \sqrt{T^2 N})$ can probably be improved to $\mathcal{O}(\rho \sqrt{TN})$, which translates into a significant performance improvement.

Perhaps more importantly, our regret bound depends on the total number of interactions. This might be improved to depend only on the number of interactions involving an honest user. The distinction is particularly important in settings where dishonest users can fabricate large numbers of interactions amongst themselves.

With plausible parameters and performance goals in the favor setting, these improvements could reduce the number of required interactions per user from tens of thousands to hundreds¹⁰ although the practical performance of the protocol might be better than the theoretical worst case, as often happens. Other performance improvements, such as distinguishing between the average and maximal payoffs or improving the hidden constant factors, would also help achieve a practical guarantee.

¹⁰

Distributed implementation: Although our protocol is tractable in the sense of polynomial-time, it is not distributed. Ideally each user’s computation and communication would be polylogarithmic in the size of the community.

The main challenge is distributing the computation of the central planner, and particularly the online learning algorithm at the core of our protocol. The current algorithm requires semidefinite programming, which must be replaced by fast distributed computation, such as random walks or sparse linear algebra. In addition to the algorithmic difficulties, the computation must be robust to manipulation by adversarial users; this might be achieved either by finding a computation which is inherently robust or by using cryptographic tools to ensure accurate computation.

Simultaneous interactions: Our protocol assumes that interactions occur one at a time: after one pair of users interacts, their reports immediately become available to the next pair of users. In most contexts this is not a realistic assumption.

When applied to a setting with k simultaneous interactions, our protocol achieves regret $\Omega(k)$. It is not hard to construct examples where the regret must necessarily be $\Omega(\log k)$. We suspect that this lower bound is tight, and that some protocol achieves regret $\mathcal{O}(\rho \log k \sqrt{TN})$. Under realistic assumptions, it may be possible to eliminate the dependence on k altogether.

Recommending vs. filtering: Our protocol answers questions of the form “should I interact with this merchant?” rather than questions of the form “which merchant should I interact with?” In many contexts the latter question is more useful—though deciding which users’ input to trust is a key difficulty for both problems. [3] provides a similar algorithm that answers “which merchant should I interact with?” but exploits the simplifying assumption that there is a single best merchant for all transactions. It is natural to try to combine these two assumptions, and to design a protocol which recommends one resource from amongst k options.

Strong signals: In our setting, most interactions provide only weak evidence about which users are honest. We might expect to be able to obtain tighter bounds when very strong signals are available. For example, a dishonest user might physically assault or threaten an honest user, resulting in a payoff of $-\rho^*$ which could never occur in an interaction between honest users. In this case, we would like to obtain a regret bound whose dependence on ρ^* was $\mathcal{O}(\rho^*N)$, independent of T .

Incentive compatibility: Our protocol has a robust theoretical guarantee, but only assuming that there are some honest users who reliably follow the protocol. It would be better if honesty were incentivized rather than assumed.

The honest players in our protocols report their payoff after each interaction, and pay a cost which depends on their reported payoff. So an honest player could always benefit by stating a lower payoff, and has no short-term incentive to play by the rules.

It is worth noting, however, that a player who follows this selfish approach will quickly find that no one wants to interact with them. Though the optimal behavior involves lying, it involves lying just enough that other users still benefit from interacting with me. In fact, a community of users lying optimally can still lead to socially optimal outcomes.

Unfortunately, in an attempt to lie optimally an honest user may inadvertently destroy value. In fact the Myerson-Satterthwaite theorem applies to our setting, even when there are just

two users,¹¹ to show that this is necessarily the case: in general it is impossible to achieve our desired guarantees at a Bayesian Nash equilibrium. We feel that our result is an attractive compromise given this impossibility result.

We could still achieve actual incentive compatibility or approximate incentive compatibility in many interesting cases where fundamental obstructions don't arise; in general, we might hope to translate positive results from the two-player iterated setting to our distributed setting.

Asymmetric interactions: Our protocols all require that interactions be symmetric *on average*, or else make use of transfers to symmetrize them. This is a very strong condition that would be nice to weaken.

One generalization is to consider sets of interactions that are balanced, so that each user has as many opportunities to give as to receive overall, even though the pairwise interactions are asymmetrical.

A more difficult generalization is to consider interactions that are inherently imbalanced, in the sense that some users systematically have more opportunities to give than to receive. In this case it is no longer possible to simultaneously maximize collective welfare for every subgroup H . We could instead try to achieve Pareto efficiency, which is a weaker condition but still meaningful if it holds for *every* set H simultaneously. For example, we might try to ensure that no group H can benefit all of its members by adopting a different internal policy and/or isolating itself from the users outside of H .

Honest supermajorities: Our guarantees are non-trivial even when the set H is a very large fraction of $[N]$. But in this case it may be able to obtain regret bounds that depend on $(N - |H|)$ rather than on N . In cases where 90% or 99% of users are honest, this could amount to a large difference in performance.

Sybilproofness: Our system bounds the damage done by the introduction of any dishonest user. In some settings it is easy for an attacker to create large numbers of dishonest users or *sybils*. Our bounds are not tight enough to discourage such an attack.

It might be possible to ensure that the introduction of a dishonest user does *no* damage to the honest users, or that the maximum possible damage done by dishonest users is bounded. For example, if transfers are available, each user could pay an upfront deposit to cover any possible damage they might inflict (in particular, a new user would need to pay a cost of at least ρ before *every* interaction with an established user, which would potentially be refunded if the interaction went well). Each honest user would need to pay a manageable cost to join the system, but the cost would be large enough that attackers can only recoup it by becoming productive members of the community.¹²

¹¹Here is a very rough sketch of the argument. Consider our setting with transfers, where $-p_0^t, p_1^t$ are independently distributed in $[0, 1]$ and are the same in every round. Because H might contain only one user or might contain both, a protocol that satisfies our guarantee must be incentive compatible and must maximize collective welfare (i.e. be ex post efficient). We can use the average size of the transfer to sell an item valued at p_0^t by the seller and p_1^t by the buyer, in a way that is incentive compatible and ex post efficient. By Myerson-Satterthwaite, we can conclude that any protocol satisfying our guarantees is not a Bayesian Nash equilibrium.

In the setting of bilateral exchange, our protocol recommends that the honest users honestly state their willingness to pay for the item, allocate it optimally, and then divide the surplus evenly. Evidently this is not a Nash equilibrium, since a player can benefit by understating their surplus. This case captures the essential and unavoidable reason that our protocol is not incentive-compatible.

¹²Impossibility results such as [16] do apply to this problem, but we are not aware of any that apply to our approach.

Composition and complex protocols: Worst-case performance guarantees allow our protocol to be used as a subroutine without compromising its correctness. Together with other building blocks, our result might be used to design more complex protocols in the same way that simple cryptographic primitives are used to build more complex cryptographic systems.